(12) UK Patent Application (19) GB (11) 2 189 112 (13) A

(21) Application No 8608698

(22) Date of filing 10 Apr 1986

(71) Applicants
STC plc

(Incorporated in United Kingdom)

10 Maltravers Street, London WC2R 3HA

(72) Inventors
Clive Ronald Carter
David Phillips

(74) Agent and/or Address for Service
S. R. Capsey, STC Patents, Edinburgh Way, Harlow, Essex CM20 2SH

(51) INT CL⁴
H04L 11/20

(52) Domestic classification (Edition I)
H4K TK

(56) Documents cited
GB A 2165422   GB A 2165125   GB A 2165124
GB A 2151880

(58) Field of search
H4K
Selected US specifications from IPC sub-class H04L

(54) Automatic telecommunication switching system

(57) An automatic telecommunication switching system handles all communications, whether voice or data, in packet switching manner. The switching network consists of a multi-stage network of co-ordinate matrix elements each of which has memory and processing means. Such an element is a single-chip VLSI device.

Each packet has a header which includes the address of the network outlet to which that packet is to be routed, each such address consisting of a digit for each stage via which the packet is to be routed. At each matrix element the appropriate digit of that address is used to route the packet to the appropriate outlet from the matrix element.
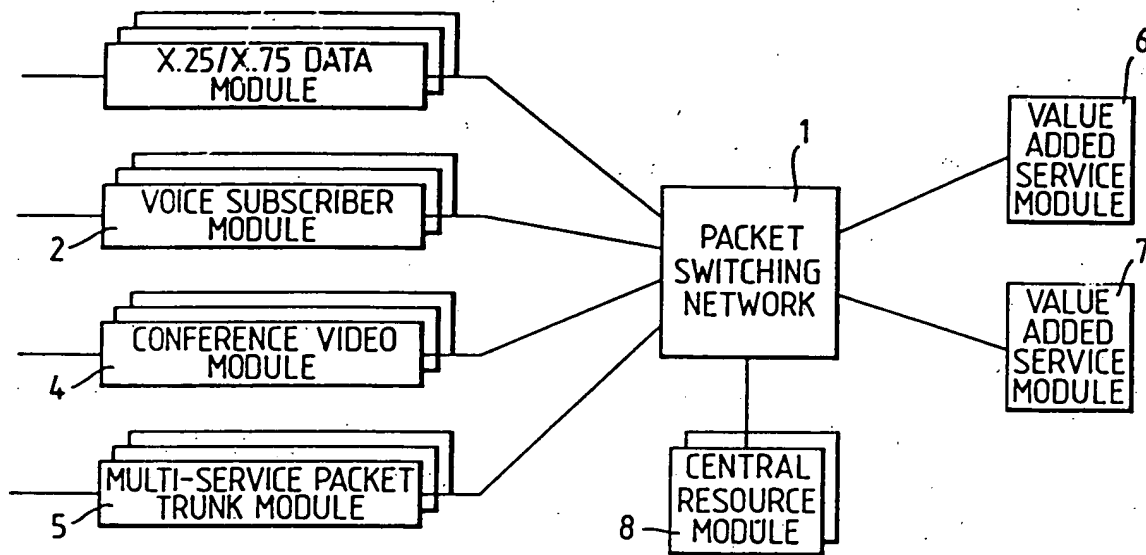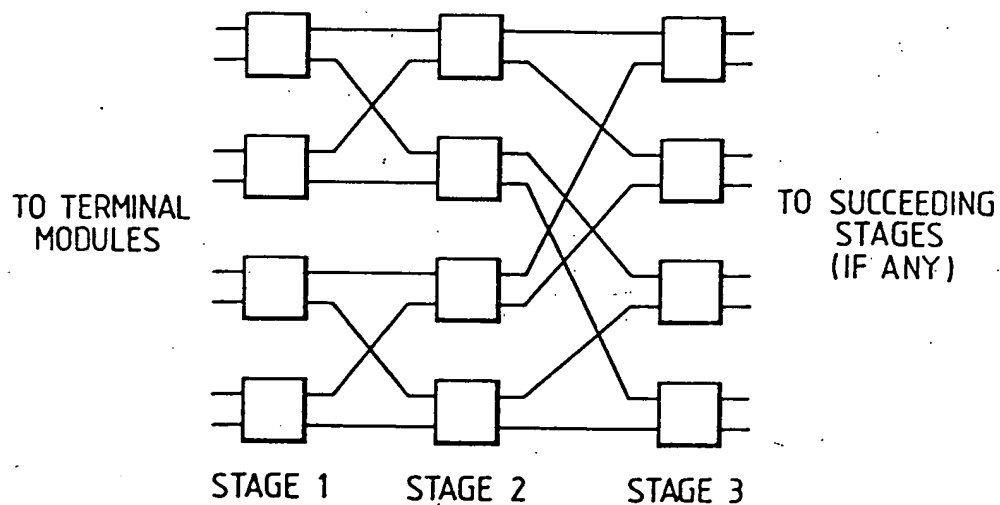
BEST AVAILABLE COPY

GB 2 189 112 A

Fig.1.

X.25/X.75 DATA MODULE

VOICE SUBSCRIBER MODULE

2

CONFERENCE VIDEO MODULE

4

MULTI-SERVICE PACKET TRUNK MODULE

5

PACKET SWITCHING NETWORK

1

VALUE ADDED SERVICE MODULE

6

VALUE ADDED SERVICE MODULE

7

CENTRAL RESOURCE MODULE

8

Fig.2.

TO TERMINAL MODULES

TO SUCCEEDING STAGES (IF ANY)

STAGE 1   STAGE 2   STAGE 3
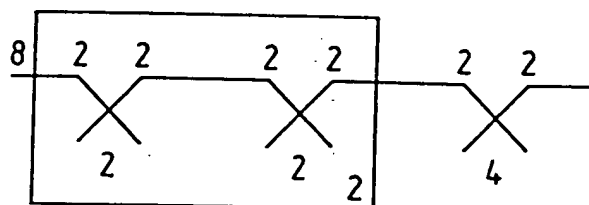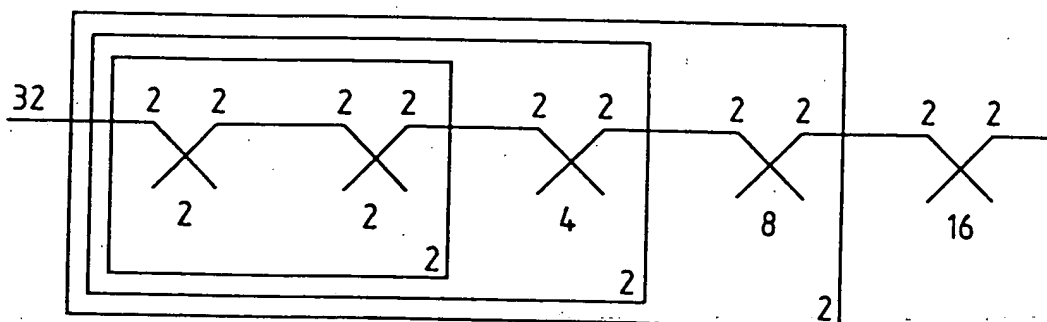
2180112

2/2

Fig.3.

Fig.4.

SPECIFICATION

Automatic telecommunication switching system

5   The present invention relates to an automatic
telecommunication switching system in which intelli-
gence including voice, is handled in packet manner.
    The emergence of new services with a wide range of
quality of service requirements presents problems for
10 current networks; existing PCM switches cannot
handle channels with bandwidths greater than
64kbits/s without major modification, and they handle
sub-rate channels inefficiently. For an integrated
switching system to handle a wide variety of data,
15 voice and visual telecommunication services it needs
a variable bit rate capability.
    The packet mode of operation provides the neces-
sary flexibility, with dynamically variable bit rate
capability both at call establishment and during a call,
20 but the current generation of packet switches provide
a bandwidth which is only sufficient to cater for data
which has a low throughput requirement and can
tolerate delay. The throughput of existing packet
switch architectures is limited by the bandwidth of
25 shared resources such as a common bus, common
buffer memory and network layer processing.
    An object of the present invention is to minimise or
overcome the effects of such limitations.
    According to the present invention there is provided
30 an automatic telecommunication switching system, in
which intelligence is handled in packet manner, in
which the system includes a plural-stage switching
network each stage of which includes a plurality of
co-ordinate matrices each of which interconnects a
35 plurality of inlets and outlets, in which each said
matrix has its own memory and processing means, so
that control of the system is at least partly distributed,
in which each intelligence packet on arriving at a said
matrix is routed into that matrix's memory means and
40 its address information examined by the matrix's
processing means to ascertain how the packet is to be
routed, and in which in response to said examination
the processing means at a matrix routes a said packet
to an outlet of that matrix appropriate to the packets'
45 address information, so that a packet is progressed
through the switching network in a stage-by-stage
manner under the control of the processing means of
the appropriate matrices.
    Embodiments of the invention will now be de-
50 scribed with reference to the accompanying drawings,
in which
    Fig. 1 is an overall "architecture" block diagram of
an automatic telecommunication switching system to
which the invention is applicable.
55   Fig. 2 is an example of a three-stage switching
network usable on the packet switching network of the
system of Fig. 1.
    Fig. 3 shows the network of Fig. 2, represented in an
alternative and simplified manner.
60   Fig. 4 is an example of a five stage matrix following
the principles of the matrix of Fig. 2, which is also
usable in the system of Fig. 1.

    The basic architecture of a multi-service packet
switch is simple and regular, consisting of a packet
65 switching network 1 connected to a variety of terminal
modules, Fig. 1. The terminal modules include port
interface modules for interfacing to external com-
munications links, subscriber 2 or trunk 3, and
exchange resource modules (subscriber information
70 database etc.). Value added service modules may also
be connected to provide such services as protocol
conversion or message handling. Thus in Fig. 1 we see
voice subscriber modules 2, X25/X75 data modules 3,
conference video modules 4, multiservice packet
75 trunk modules 5 and value-added modules 6, 7.
    The switching network consists of a regular array of
identical packet switching elements organised as a
folded network matrix. That is, all terminals to be
interconnected are connected to one side of the
80 matrix, and a connection is set up by going into the
matrix and then back out of it. The packet switching
elements have the processing power and memory
needed to execute the algorithms to route packets
between the terminal modules. The number of switch-
85 ing stages required depends on the number of
terminals, but because of the folded network arrange-
ment the switch size can be increased by adding
additional switching stages without requiring any
rearrangement of the existing network or terminal
90 connections. Only a single functional unit, the packet
switching element, is required to build uo the entire
switching network. This element is in effect an
electronic analogue of a cross-bar switch, with some
of its own processing.
95   Terminal modules are connected to the packet
switching network via a standard serial interface, so
that new terminal types can be added without
affecting the packet switching network.
    A packet "datagram" store and forward switching
100 technique is used whereby each datagram packet
carries a routing address in its header, and by
reference to this address the switching elements can
forward the datagrams over the appropriate path.
    The packet switching network overcomes many of
105 the drawbacks of existing packet switches which rely
on shared buses for packet transfer, and central
processors for call control; the range of application of
such systems and their ability to grow is seriously
limited. In contrast the packet switching network used
110 in the present system consists of an array of identical
switching elements connected together in a matrix.
Thus the traffic is distributed across the matrix and as
the exchange grows the matrix also grows maintain-
ing a constant traffic level on the links. An example of a
115 three stage matrix of switching elements with four
connections each is shown in Fig. 2, and Fig. 3 shows
3-stage and 5-stage matrix examples in an alternative
shorthand representation. Note that in these networks
each switching element is a 2 x 2 co-ordinate switch.
120 The use of a folded single-sided structure allows
continuous expansion by the addition of further
switching stages as the number of terminals increases
without the need to change connections, or to
predetermine the maximum size of a switch.

The packet switching element is a single VLSI circuit incorporating *n* identical full duplex links plus internal memory and logic for packet buffering and routing. The optimum value for the number of links (n) lies
5 between 4 (the minimum practical) and 16. The speed of the links must be of the order of tens of megabits per second to give the wide bandwidth and the short delays required to support a range of services.

Each link carries an incoming and an outgoing serial
10 bit stream, which operates asynchronously eliminating the need for synchronous timing signals and their distribution throughout the switch network. Packets received on the incoming links are autonomously written into buffer space in RAM (Random Access
15 Memory) by DMA (Direct Memory Access). Then an algorithm is executed by the internal processing logic of the switching element to route the received packet to one of the outgoing links according to a routing code in a header field of the packet. Packets are
20 transmitted on the outgoing links from buffer RAM by DMA.

The maximum packet size must be set large enough to keep the processing overheads small in proportion to the transmission time, but small enough such that
25 the buffer space required does not exceed the available memory size. Packet size may differ in different systems, but one size usable is 128 bytes.

Commercially available components for use in a laboratory demonstration model provide 4 links at 10
30 or 20 Mbit/s, 4 kilobytes of RAM and have a machine cycle time of 50 nanoseconds. Such a size is adequate for a laboratory model provided that the routing algorithms are simple to mimimise processing overheads and that a reasonable compromise packet size
35 can be established. If these parameters are unsuitable for a product, custom VLSI can be developed.

Port interface modules support the external ports of the switch and provide the service required for particular types of traffic. These modules are each
40 connected to the network via a standard asynchronous serial interface. A variety of port interface module types are required to support various port interfaces. Some examples of the types of port interfaces which may be supported are:
45 (a) X.25 packet data subscriber port
(b) Voice subscriber port
(c) 2Mbit/s conference video circuit port
(d) Multi-service packet data trunk port.

Each module incorporates dedicated interface cir-
50 cuits to support the appropriate external ports and a processor plus memory to provide packet buffering, external protocol handling, call processing, and transfer of data and signalling across the switch network. Sufficient processor power is provided in the modules
55 to support all call processing functions.

The port interface modules are not described herein in detail, but their functions include:
(i) send and receive data (including voice where appropriate) and signalling over the external ports
60 (ii) assemble/disassemble data packets if the external information stream is not packetised
(iii) fragment packets which are too large for transfer across the switching network
(iv) call establishment in co-operation with the module
65 handling the other call half and centralised resources

(v) send and receive data and signalling over the internal switch network providing reliable transfer by the use of error detection or correction code and retransmission where required.

Note that since we are using a folded network, each call is set up in two halves, one into the network, and one out of it.

Some of the requirements to support particular services are as follows:
75 (a) X.25/X.75 packet data service requires reliable sequenced data transfer, which requires the detection of lost or corrupted packets and their retransmission. Also packet sizes of up to 4 kilobytes may be requested at call establishment, and this requires packet frag-
80 mentation before transfer across the switch network.

(b) Voice service is tolerant of loss, but sensitive to delays, so it is usual to accept packet loss without retransmission and to keep voice packets short to minimise packet assembly delays. Apart from the
85 widely used 64 kbit/s PCM, a number of other voice coding schemes are proposed for public networks providing lower bit rates or higher quality. Voice processing may also be required for silence suppression or echo cancellation.
90 (c) A trunk port interface module should support one trunk circuit operating at an appropriate speed of the national hierarchy, e.g. 2 Mbit/s or 8 Mbit/s. Because the terminal modules operate through a standard interface, new terminal types can be added
95 without affecting the existing packet switching network.

The peripheral modules, where subscriber lines and trunks are terminated, incorporate microprocessors which provide packet buffering, external protocol
100 handling, transfer of data and signalling across the switch network, and many of the call-processing functions. Such an arrangement is not only more efficient than a fully centralised system, but allows the processing capacity to increase in proportion to the
105 number of terminations supported.

Not all switching functions, however, are best provided by a distributed architecture. Certain functions such as system maintenance, control of external links to remote operations and maintenance centres,
110 bulk data storage and subscriber and routing information are best provided by centralised resources, as seen from the inclusion in Fig. 1 of a central resource module 8.

The processors communicate with each other by
115 transferring message packets across the same switch network as the data (intelligence) packets. The call processing functions of the central resource processors 7 consist mainly of responding to requests from the peripheral module processors for address or routing translation and determination of terminating line or trunk equipment number and selection of a trunk in a trunk group where appropriate, Given the destination equipment number the derivation of a routing address to route packets from source to destination peripheral modules is simple and the peripheral modules can then exchange messages to effect call establishment. Where however, the switch is being used in a leased line or other slave application the central resource processor may itself initiate path establishment.

The internal protocol used in the system is the protocol used to exchange data and signalling messages between terminal modules (port interface modules and resource modules) across the network.

5 The packet format includes a header which has two fixed length fields, a routing address field and a packet length field. The following fields, namely the information field and the error check code, are forwarded transparently by the switching elements. That is, they

10 are not altered in any way by those elements.

The routing address consists of a number of routing 'digits', where one 'digit' determines the routing through one switching element. If the switching elements are 2 x 2 switches, each such digit can have

15 one bit, 1 or 0 dependent on which outlet is needed. In the case of a 16 x 16 switching element, each such digit has four bits. The number of routing digits required to select a path through the matrix depends on the number of stages to be traversed and is thus variable,

20 but the routing digits are padded out to maintain a fixed length field to minimise processing overheads. The routing address required to select a path between a pair of modules is determined by the source module at virtual call establishment and this address is used

25 on all packets for the duration of the call, thus ensuring that all packets of a call follow the same path and therefore are delivered in sequence.

The length field contains a binary count of the length of the information field in bytes, m and the error check

30 code protects against errors in link transmission or in the switching elements. As the links are short and situated within the "benign" environment of an exchange transmission errors should be negligible under normal conditions. The error check code is still

35 required however, to detect errors under hardware failure conditions. Link error detection has no advantage with the low error rate expected, so error detection is performed end to end across the network.

A network such as that of Fig. 2, is in effect a

40 "multi-rooted tree" in which switching element has the same number of connections to "parents" as "children". If all packets travelled all the way up to a root element there would be no concentration of packets onto upper links. However, since packets need

45 only travel up as far as a common "ancestor", successive upper levels are in fact progressively under used. Thus the network can be extended in levels indefinitely without overloading.

For a multi-rooted tree such as this the choice of

50 "parent" link on the upward journey is arbitrary and can be dynamically chosen to avoid congestion. The upward path is used to get to a common ancestor, and any common ancestor will do. Traversing these arbitrary upward links constitutes an inefficiency of

55 50% as half the links traversed do not switch the packet toward its destination. The round trip delay is therefore twice that necessary.

The Inmos Transputer is a fast single chip computer designed to interconnect with other transputers in a

60 distributed multi-transputer architecture. The language in which it is programmed, called "Occam", is designed to support communications between concurrent processes distributed across interconnected Transputers. The design goals behind the Transputers

65 are thus compatible those of a switching system with distributed functionality. In particular the integration of communications (four 10 to 20 Mbit/sec. links), storage (4 Kbytes) and processing (10 MIPS) onto a single chip makes the transputer a suitable compo-

70 nent for building a packet switching network.

For the hardware engineer the transputer promises to simplify the design by reducing the amount of hardware required external to the transputers themselves. Any two communicating transputers can be

75 directly connected via a two-wire (plus ground) communication link, yet transputers do not need to be synchronous to communicate as they can operate reliably from independent 5 MHz clocks. The point-to-point (transputer-to-transputer) communication links

80 operate serially at 10 or 20 Mbits/sec. and require no buffering over short distances. These features make large systems consisting of many transputer switching elements feasible since neither data nor a synchronising clock needs to be bussed around the entire

85 system.

The current version of the transputer has four ports, each with an incoming path and an output path. Hence it is usable as a 2 x 2 co-ordinate switch.

For a packet switch application, the main limitation

90 of the transputer is the small number of other transputers it can directly communicate with. Structures requiring connectivity of more than four would require each node to consist of more than one transputer. However, it is readily possible to develop a

95 switching element of larger connectivity, e.g. the 16 x 16 element mentioned briefly above.

CLAIMS

1. An automatic telecommunication switching system, in which intelligence is handled in packet

100 manner, in which the system includes a plural-stage switching network each stage of which includes a plurality of co-ordinate matrices each of which interconnects a plurality of inlets and outlets, in which each said matrix has its own memory and processing

105 means, so that control of the system is at least partly distributed, in which each intelligence packet on arriving at a said matrix is routed into that matrix's memory means and its address information examined by the matrix's processing means to ascertain how the

110 packet is to be routed, and in which in response to said examination the processing means at a matrix routes a said packet to an outlet of that matrix appropriate to the packets' address information, so that a packet is progressed through the switching network in a

115 stage-by-stage manner under the control of the processing means of the appropriate matrices.

2. A system a claimed in claim 1, and in which the connections of subscriber's lines, trunks and facilities circuits to the network are each made via interface

120 modules which, at least in the cases of the lines and trunks, include processing means.

3. A system as claimed in claim 1 or 2, and in which some switching functions, e.g. system maintenance, are provided by centralised control means.

125 4. A system as claimed in claim 1, 2 or 3, and in which interprocessor communications between module interfaces and/or the processing means at the matrices is effected over the same physical path as that used for the intelligence.

130 5. A system as claimed in claim 1, 2, 3 or 4, in which

for a connection to be set up the number of stages that
connection extends into the network is variable, the
shortest possible path being used.

 6.   A system a claimed in claim 1, 2, 3, 4 or 5 and in
5 which each said switching element is a single chip
VLSI semiconductor device.

 7.   An automatic telecommunication switching
system, substantially as described with reference to
the accompanying drawings.

10

New claims or amendments to claims filed on 28th
July, 1986.

 8.   An automatic telecommunication switching
15 system in which intelligence is handled in packet
manner, in which the system includes a plural-stage
folded switching network with the terminals served by
the system connected to one side of the network, in
which each connection between two of said terminals
20 is set up via two half-connections one of which
extends into the network from said one side and the
other of which comes out of the network to said one
side, in which each stage of the switching network
includes a number of co-ordinated matrices each
25 including a number of semiconductor cross-points,
each such matrix interconnecting a plurality of matrix
inlets and matrix outlets, in which each said switching
matrix has its own memory means and its own
processing means so that control of the system is at
30 least partly distributed, in which when an intelligence
packet traversing the network between two of said
terminals arrives at a said matrix that packet is routed
into that matrix's memory means and its address
information is examined by that matrix's processing
35 means to ascertain how the packet is to be routed, and
in which in response to said examination the proces-
sing means at a said matrix routes a said packet to an
outlet of that matrix appropriate to a part of the
address information of that packet, so that a packet is
40 progressed through the switching network in a
stage-by-stage manner under the control of the
processing means of the appropriate ones of said
matrices.

 9.   A system as claimed in claim 8, in which for a
45 connecion to be set up the number of stages that
connection extends into the network is variable, the
shortest possible path being used.

 10.   A system as claimed in claim 8 or 9, in which
the switching network is a three-stage network
50 arranged in the way shown in Figs. 2 and 3.

 11.   A system as claimed in claim 8 or 9, in which
the switching network is a five-stage network
arranged in the way shown in Fig. 4.

 12.   A system as claimed in claim 8, 9, 10 and 11, in
55 which the terminals served by the system, which
includes subscribers' lines, trunks and facilities cir-
cuits, are each connected to the network via an
interface module which, at least in the cases of the
lines and trunks, includes processing means.

60 13.   A system as claimed in claim 8, 9, 10, 11 or 12,
in which inter-processor communication between the
processing means at the matrices and/or the proces-
sing means in the interface modules is effected over
the same physical path as that used for the intelli-
65 gence.